# Robust Massively Parallel Sorting

Michael Axtmann, Peter Sanders
*Karlsruhe Institute of Technology*, *Karlsruhe, Germany*
*{michael.axtmann,sanders}@kit.edu*

arXiv:1606.08766v3 [cs.DC] 2 Nov 2016

***Abstract*—We investigate distributed memory parallel sorting algorithms that scale to the largest available machines and are robust with respect to input size and distribution of the input elements. The main outcome is that four sorting algorithms cover the entire range of possible input sizes. For three algorithms we devise new low overhead mechanisms to make them robust with respect to duplicate keys and skewed input distributions. One of these, designed for medium sized inputs, is a new variant of quicksort with fast high-quality pivot selection.**

**At the same time asymptotic analysis provides performance guarantees and guides the selection and configuration of the algorithms. We validate these hypotheses using extensive experiments on 7 algorithms, 10 input distributions, up to 262 144 cores, and varying input sizes over 9 orders of magnitude. For "difficult" input distributions, our algorithms are the only ones that work at all. For all but the largest input sizes, we are the first to perform experiments on such large machines at all and our algorithms significantly outperform the ones one would conventionally have considered.**

## I. Introduction

Sorting is one of the most fundamental and widely used non-numerical algorithms. It is used to build index data structures, to establish invariants for further processing, to bring together "similar" data, and in many other applications. This wide variety of applications means that we need fast sorting algorithms for a wide spectrum of inputs with respect to data type, input size, and distribution of keys. The motivation for this paper is that parallel algorithms currently do not cover this spectrum of possible inputs for very large parallel machines. Although hundreds of papers on parallel sorting have been published, there is only a small number of practical studies that consider the largest machines with many thousands of processors (PEs). The few studies known to us mostly concentrate on very large random inputs. Most of these algorithms become slow (or simply crash) when applied to worst case inputs where the location of the input data is correlated with key values (skewed inputs) or which may have large numbers of duplicated keys.

Even for random inputs, these algorithms are slow on small inputs, where their running time is dominated by a large number of message startup overheads. Note that sorting small inputs becomes important for the performance of parallel applications when it is needed in some frequently repeated coordination step between the PEs. For example, many applications perform load (re)balancing by mapping objects to space filling curves and sorting them with respect to this ordering [1]. The scalability of the sorting algorithm may then become the limiting factor for the number of time steps we can do per second. Note that in this case it is of secondary importance whether the sorting algorithm itself is efficient compared to a sequential one – what matters is that it is faster than the local computations between rebalancing steps. Siebert and Wolf [2] give another extreme example, pointing out that the `MPI_Comm_Split` operation requires sorting with exactly one input per PE.

The subject of this paper is to systematically explore the design space of parallel algorithms for massively parallel machines and propose robust algorithms for the entire spectrum of possible inputs. As a first design decision, we restrict ourselves to comparison-based algorithms because they work for arbitrary data types and are usually less dependent on the distribution of key values than non-comparison based algorithms like radix sort. The disadvantage of more local work for comparison-based algorithms is often irrelevant for the largest machines since communication overhead is the main bottleneck. For similar reasons, we abstain from low-level tuning like exploiting SIMD-instructions or using hybrid codes with node-local shared memory algorithms.

Let $n$ denote the input size and $p$ the number of PEs. We assume that each PE has a local input of size $\mathcal{O}(\frac{n}{p})$, i.e, the input is "reasonably" balanced. We consider three major issues with respect to the robustness of a massively parallel sorting algorithm: Its scalability, i.e., its running time as a function of $n$ and $p$, how the algorithm behaves with respect to skewed input distributions and how it handles repeatedly occurring keys.

*Our Contributions.* A methodological contribution is that we address the scalability issue by reconciling theory and practice (to some extent). Classical theoretical algorithms concentrate on the issue how fast sorting algorithms can be made for small inputs. For example, Cole's celebrated parallel mergesort [3] runs in time $\mathcal{O}(\log p)$ for $n = p$ on a PRAM. However, this result and the PRAM machine model are considered quite impractical. We show that simpler algorithms with polylogarithmic running time are actually practical for small $n/p$. Using asymptotic analysis in a more realistic yet simple complexity model that distinguishes between startup latencies and communication bandwidth, we get a consistent ranking of a large number of algorithms with respect to their performance over the entire range of values for $n/p$. We propose three robust sorting algorithms for different input sizes which can be used to span the entire parameter space. The first sorts while data is routed to a single PE, the second is a simple yet fast work-inefficient algorithm with logarithmic latency (RFIS), the third is an efficient variant of parallel Quicksort with latency $\mathcal{O}(\log^2 p)$ (RQuick) and the fourth has latency $\mathcal{O}(p^{1/d})$ when we allow data to be moved $d$ times (RAMS).

RFIS and RAMS improve our own results from a mostly theoretical paper [4] that includes a nonrobust proof of concept implementation and experiments on up to $2^{15}$ PEs. However,

these experiments only used uniform inputs and had highly fluctuating running times. We also concentrated on rather large inputs (RFIS is used only as a nonrobust subroutine for ranking samples). Our contribution to RFIS and RAMS are robust implementations and extensive experiments with various input distributions on up to $2^{18}$ PEs. We use the input distributions from Helman et al. [5], who pioneered robustness experiments on parallel sorting algorithms. Our quicksort algorithm is new, incorporating a fast and accurate splitter selection algorithm and low overhead measures to handle skew and repeated keys. We conjecture that the expected asymptotic running time for worst case inputs (skewed but with unique keys) is the same as for best case inputs where local data is split perfectly in each step. This algorithm closes the gap between very small and large inputs where RAMS is still inefficient but the local work of RFIS begins to dominate its startup latencies. Finally, we compare our algorithms directly with several state of the art implementations considered in a recent study [6]. It turns out that sorting while gathering data on a single PE "sorts" very sparse inputs the fastest. RFIS is the fastest algorithm on sparse and very small inputs. Our new quicksort outperforms its competitors on small inputs, in the range of $2^3$ to $2^{14}$ elements per core, robustly for all input instances. HykSort [6] is slightly faster than RAMS for large random inputs but less robust. Classical sample sort and related algorithms that communicate the data only once are very slow even for rather large $n/p$ [5], [7]–[10]. Classical bitonic sorting [11], [12] is somewhat competitive only for a rather narrow range of input sizes.

*Paper Overview.* We give preliminary definitions in Section II and describe building blocks for random shuffling and median estimation in Section III. Section IV compares eleven parallel sorting algorithms analytically and qualitatively. Our main new algorithms are described in more detail in Sections V–VI. Section VII describes an extensive experimental evaluation.

## II. Preliminaries

The input of sorting algorithms are $n$ elements with $\mathcal{O}(n/p)$ elements on each PE. The output must be globally sorted, i.e., each PE has elements with consecutive ranks. We also want $\mathcal{O}(n/p)$ output elements on each PE. Sometimes we more concretely consider perfectly balanced inputs and an output with at most $(1+\epsilon)n/p$ elements per PE for some small positive constant $\epsilon$.

A naive approach to generating unique keys appends unique identifiers to the elements. Then, lexicographic ordering makes the keys unique. However, this approach makes sorting considerably more expensive due to higher communication volume and more expensive comparisons.

*Model of Computation.* A common abstraction of communication in supercomputers is the (symmetric) single-ported message passing model. Algorithms in this model are often implemented using the MPI, a standardized message passing interface [13], that supports point-to-point message exchange as well as a wide variety of collective operations like broadcast, reduction, prefix sum, etc. It takes time $\alpha + l\beta$ to send a message of size $l$ machine words. The parameter $\alpha$ defines the startup overhead of the communication, while $\beta$ defines the time to communicate a machine word. Appendix A gives further details of the model of computation.

*Hypercube Algorithms.* A hypercube network of dimension $d$ consists of $p = 2^d$ PEs numbered $\{0, \ldots, p-1\}$. Two nodes $a$ and $b$ are connected along dimension $i$ if $a = b \oplus 2^i$. For this paper, hypercubes are not primarily important as an actual network architecture. Rather, we extensively use communication in a conceptual hypercube as a design pattern for algorithms. More specifically the hypercube algorithms shown in Algorithm 1 iterate through the dimension of the hypercube. Depending on how this template is instantiated, one achieves a large spectrum of global effects. To describe and understand hypercube algorithms, we need the concept of a subcube. A $j$-dimensional subcube consists of those PEs whose numbers have the same bits $j..d-1$ in their binary representation. Basic operations such as *all-reduce*, *all-gather*, and routing data for random start or destination nodes [14] need only $\mathcal{O}(\alpha \log p)$ startup overheads overall. We call the hypercube algorithm which places all elements on all PEs in sorted order *all-gather-merge*. This operation runs in time $\mathcal{O}(\beta p|a| + \alpha \log p)$. Appendix B describes these operations in more detail.

**Algorithm 1** Hypercube algorithm design pattern

```
Computation on PE i
for j:= d − 1 downto 0 do                    ▷ or 0..d − 1
    send some message m to PE i ⊕ 2^j
    receive message m′ from PE i ⊕ 2^j
    perform some local computation using m and m′
```

## III. Building Blocks

### A. Randomized Shuffling on Hypercubes

It is folklore (e.g. [15]) that data skew can be removed by shuffling the input data randomly. Helman et al. [5] implemented this for the sample sort algorithm by directly sending each element to a random destination. Note that this incurs an overhead of about $\alpha p + \beta n/p$. In Appendix C we propose a technique for small inputs which runs in $\mathcal{O}((\alpha + \beta\frac{n}{p}) \log p)$.

### B. Approximate Median Selection with a Single Reduction

Siebert and Wolf [2] consider the case where $n = p$ and propose to select splitters for parallel quicksort using a ternary tree whose leaves are the input elements. At each internal node, the median of the children elements is passed upward. Dean et al. [16] show that for randomly permuted inputs and a balanced tree this is a good approximation of the median, i.e., with high probability, one gets rank $n/2(1 \pm 2n^{-0.37})$. Siebert and Wolf [2] do not permute randomly. Even when $p$ is a power of two, their method does not produce a completely balanced tree so that their implementation remains a heuristic.

We fix these restrictions by using a binary tree and working with randomly permuted inputs: Consider a tuning parameter $k$ (even). Assume that the local data is a sorted sequence

$a[1..m]$ and that $m$ is even. Each PE is a leaf of the binary tree and forwards $a[m/2-k/2+1..m/2+k/2]$ up the tree – a local approximation of the elements closest to the median. Undefined array entries to the left are treated like a very small key and undefined entries to the right as very large keys. If $m$ is odd, we flip a coin whether $a[\lfloor m/2\rfloor-k/2+1..\lfloor m/2\rfloor+k/2]$ or $a[\lceil m/2\rceil-k/2+1..\lceil m/2\rceil+k/2]$ is forwarded. Internal nodes merge the received sequences and use the result as their sequence $a$ analogously to the leaves. At the root, we flip a coin whether $a[k/2]$ or $a[k/2+1]$ is returned. Note that in most MPI implementations, this algorithm can be implemented by defining an appropriate reduction operator. The overall running time is $\mathcal{O}(\alpha \log p)$. For randomly permuted local inputs of size $n/p$, it is easy to show that our scheme yields a truthful estimator for the median, i.e., we get a result with expected rank $n/2$. We conjecture that similar quality bounds as from Dean et al. [16] can be shown for our scheme, i.e., that we get rank $n/2(1 \pm cn^{-\gamma})$ with high probability for some constant $c$ and $\gamma$. Our experiments in Appendix H indicate that the approximation guarantee of the binary tree is about $1.44n^{-0.39}$ whereas the approximation guarantee of the ternary tree is about $2n^{-0.37}$.

## IV. Sorting Algorithms from Tiny to Huge Inputs

We now outline a spectrum of parallel sorting algorithms (old and new), analyze them and use the result to compare their performance for different values of $n$ and $p$. Table I summarizes the results, roughly going from algorithms for small inputs to ones for larger and larger inputs. We only give the cost terms with respect to latency ($\alpha$) and communication volume ($\beta$) per PE because the relevant local work is always the same $\mathcal{O}(\frac{n}{p}\log n)$. We now compare the algorithms with regard to running time and robustness. Appendices D and E describe the algorithms in more detail.

We recently proposed Fast Work-Inefficient Ranking (FIR) [4], a simple but fast ranking algorithm for small inputs where the PEs are arranged in an array of size $\mathcal{O}(\sqrt{p}) \times \mathcal{O}(\sqrt{p})$. The result of the algorithm is that each PE knows the global rank of all input elements in its row. In theory this is a very good algorithm for $n = \mathcal{O}(\sqrt{p})$ since we have only logarithmic delay. In Section V we propose Robust fast work-inefficient sort (RFIS) which makes the algorithm robust against identical keys and converts its output to a classical sorted permutation of the input.

TABLE I: Complexity of various parallel sorting algorithms. Implicit $\mathcal{O}(\cdot)$.

| Algorithm | Latency $[\alpha]$ | Comm. Vol. $[\beta]$ | Remarks |
|---|---|---|---|
| (All-)gather-merge | $\log p$ | $n$ | II |
| RFIS [4] | $\log p$ | $n/\sqrt{p}$ | here:robust |
| Bitonic [12] | $\log^2 p$ | $\frac{n}{p}\log^2 p$ | deterministic |
| Minisort [2] | $\log^2 p$ | $\log^2 p$ | $n = p$ |
| HC quicksort [17] | $\log^2 p$ | $\frac{n}{p}\log p$ | best case |
| +median of medians [18] | $\log^2 p$ | $p + \frac{n}{p}\log p$ | average case |
| HC quicksort Sect. VI | $\log^2 p$ | $\frac{n}{p}\log p$ | here:robust |
| HykSort [6] | $\geq k \log_k p$ | $\geq \frac{n}{p}\log_k p$ | not robust |
| AMS [4] | $k \log_k p$ | $\frac{n}{p}\log_k p$ | here:robust |
| Sample sort [7] | $\geq p$ | $\geq n/p$ | +sampling cost |
| Multiway merges. [4], [8] | $\geq p$ | $\geq n/p$ | |

For small inputs the running time of Bitonic sort [11], [12] is dominated by $\log^2 p$ startup overheads. This gives it a similar running time as the parallel quicksort algorithms to be discussed next. However, for $n = \omega(p\alpha/\beta)$ the term $\beta\frac{n}{p}\log^2 p$ dominates – all the data is exchanged $\log^2 p$ times. This makes it unattractive for large inputs compared to quicksort and other algorithms designed for large inputs.

There are many parallel variants of quicksort. Hypercube quicksort uses the hypercube communication pattern, is simple and can exploit locality of the network. In the simplest original form by Wagar [17], PE 0 uses its local median as a pivot. This is not robust against skew and even for average case inputs it only works for rather large $n$. One can make this more robust for average case inputs by using the global median of the local medians [18]. However, this introduces an additional $\beta p$ term and leads to non-polylogarithmic execution time. Therefore in Section VI we make quicksort robust against skew and propose to use a median selection algorithm that is both fast and accurate. Non-hypercube distributed memory quicksort leads to more irregular communication patterns and has its own load balancing problems because we cannot divide PEs fractionally. Siebert and Wolf [2] exploit the special case $n = p$ where this is no problem.

Quicksort exchanges the data at least $\log p$ times. More general approaches partition the input into $k$ parts on each recursion level. Those algorithms move the data $\mathcal{O}(\log_k p)$ times and the lower bound of the running time is

$$\Omega\left(\frac{n}{p}\log n\right) + \beta\frac{n}{p}\log_k p + \alpha k \log_k p \ .$$

Many algorithms have been devised for the special case $k = p$ [7], [19]. These single level algorithms exchange data only once. Sample sort [7] achieves the desired bound for relatively large inputs of $n = \Omega(p^2/\log p)$. Helman et al. [5] propose initial random shuffling to make samplesort robust against skewed input distributions (this was already proposed in [20]). One can even achieve perfect partitioning by finding optimal splitters in a multiway mergesort approach [4], [8], [10]. This requires $n = \Omega(p^2 \log p)$ for achieving the desired bound. Gerbessiotis and Valiant [9] develop a multilevel sample sort for the BSP model [21]. However, implementing the data exchange primitive of BSP requires $p$ startup overheads in each round. Cole and Ramachandran [22] and Blelloch et al. [23] give multi-level sorting algorithms for cache oblivious shared memory models. Translating them to a distributed memory model appears nontrivial and is likely to result in worse bounds.

Sundar et al. [6] develop and implement *HykSort*, a multi-level generalization of hypercube quicksort. HykSort is "almost" robust against skew. However, in the worst case, data may become severely imbalanced. The paper mentions measures for solving this problem, but does not analyze them. Furthermore, HykSort uses the operations `MPI_Comm_Split` whose current implementations need time $\Omega(\beta p)$. This is why we put a "$\geq$" in the corresponding row of Table I. HykSort is also not robust with respect to duplicate keys.

Our recently published $k$-way sample sort algorithm, *AMS-sort* [4], guarantees a maximum imbalance of $\epsilon\frac{n}{p}$ for arbitrary inputs. AMS-sort avoids the data imbalances within target subcubes that may make HykSort inefficient. Our paper [4] describes (but does not implement) a way to select communication partners such that each PE sends and receives $\Theta(k)$ messages for a $k$-way data exchange. Similarly, a tie-breaking technique is proposed (but not implemented) which simulates unique keys with very low overhead. In Appendix G we describe RAMS, our robust implementation of AMS-sort (see also Section V). Note that the running time of the multi-level algorithms described above with $\Theta(\log p)$ levels is $\Omega(\log^2 p)$. For example, the running time of AMS-sort is then $\mathcal{O}(\log^2 p \log\log p)$. However, in this case quicksort is the simpler algorithm, has a startup latency of $\mathcal{O}(\log^2 p)$ and smaller constant startup latency factors.

## V. Robust Fast Work-Inefficient Sorting and Multi-Level Sample Sort

We propose a fast sorting algorithm (RFIS) for sparse and very small inputs that is robust against duplicate keys. Our algorithm computes the rank of each element in the global data. Then, optionally, a data delivery routine sends the elements to the appropriate PEs according to their ranks. For ranking the input data, we use the algorithm we recently proposed [4] (see also [24] and Appendix D1). Unfortunately the proposed approach calculates the same rank for equal elements. However, our data delivery technique requires unique ranks in the range of $0..n-1$ as input to achieve balanced output. Consider a tie-breaking policy that logically treats an element as a quadruple $(x, r, c, i)$ where $x$ is the key, $r$ the row, $c$ the column, and $i$ the position in the locally sorted data. We can then define a total ordering between these quadruples as their lexicographical order. We now exploit the fact that we do not rely on the true column $c$ (row $r$) but on the information weather an element came from the left (above) or from the right (below). A modified all-gather-merge and a modified compare function implement that policy without any communication overhead for communicating the $(r, c, i)$ information. Appendix F describes this policy in more detail. After computing the ranks of the input elements, we have to redistribute them to get sorted output. The element with rank $i$ is mapped to PE $ip/n$. We now exploit the fact that each PE column stores the complete ranked input. We thus can afford to discard all elements that are not mapped to the same column without losing any element. The data delivery problem now can be solved locally in each column. We use a hypercube algorithm for routing the elements to their correct row. In the $i$-th iteration of this algorithm, at most $n/(p^{\frac{1}{2}}2^i)$ elements have to be delivered in their respective subcube. Even if all these elements are concentrated on the same PE, this requires at most $\alpha + \beta n/(p^{\frac{1}{2}}2^i)$ time. Summing this over all iterations yields overall time $\mathcal{O}(\alpha \log p + \beta\frac{n}{\sqrt{p}})$, i.e., the asymptotical running time is the same for just ranking or ranking plus data delivery.

The prototypical implementation of AMS in [4] does not ensure that each PE communicates with at most $\mathcal{O}(k)$ PEs in each level of recursion (deterministic message assignment). Similarly, [4] proposes but does not implement a tie-breaking scheme to become robust against duplicate keys with almost no communication overhead. in Appendix G, we present a complete implementation of AMS-sort (RAMS) including both improvements outlined above.

**Algorithm 2** Robust Quicksort on Hypercubes

**Input:** $a = \{a_1, \ldots, a_{n/p}\}$ a set of local input elements, $p = 2^d$ number of PEs, PE number $i$
$a \leftarrow$ randomly redistribute $a$ ▷ see Section III-A
SORT($a$)
**for** $j$:= $d-1$ **downto** $0$ **do** ▷ iterate over cube dims
  $s \leftarrow$ calculate splitter in parallel ▷ see Section III-B
  **if** ISEMPTY($s$) **then return** $a$ ▷ no elements in cube
  split $a$ into $L \cdot R$ ▷ see text
  $i'$:= $i \oplus 2^j$ ▷ communication partner
  **if** $i' < i$ **then**
    send $L$ to $i'$ and receive $R'$ from $i'$
    $a \leftarrow$ merge $R$ with $R'$
  **else**
    send $R$ to $i'$ and receive $L'$ from $i'$
    $a \leftarrow$ merge $L$ with $L'$
**return** $a$

## VI. Robust Quicksort on Hypercubes

Previous implementations of hypercube-based quicksort (e.g. [17], [18]) have three major drawbacks. First, they are not robust against duplicates as no tie-breaking is performed. Second, the algorithms do not consider non-uniform input distributions. It can be shown (e.g. [14]) that skewed inputs exist where the imbalance increases rapidly over the first $\frac{\log p}{2}$ recursions. After $\frac{\log p}{2}$ recursions, $\lceil\sqrt{p}\rceil$ PEs hold $\lfloor n/\sqrt{p}\rfloor$ elements each. This bound even applies for unique input keys and true medians as a global splitter. Third, a fast high quality splitter selection with median approximation guarantees is essential. Otherwise, quicksort gets impractical for large $p$ as the load imbalance accumulates or the splitter selection dominates the running time for small inputs. Previous quicksort implementations use one of two different approaches to calculate the splitter. The first approach [17] selects a random PE which just broadcasts its local median. The second approach [18] gathers the local median of each PE and calculates the median of medians.

We propose a simple but robust hypercube quicksort that overcomes these problems. Algorithm 2 gives pseudocode. First, our algorithm randomly redistributes the input (see Section III-A). Then the data is sorted locally. The main loop goes through the dimensions of the hypercube, highest dimension first. In iteration $j$, the fast median selection algorithm from Section III-B is used to approximate the median of the data in each $j+1$-dimensional subcube. This key is used as a splitter $s$ for that subcube. For robustness against repeated keys, we use a low overhead tie-breaking scheme that works locally without the need to communicate additional information: Suppose a PE holds a sorted sequence $a$ of the form $a = a_\ell \cdot s^m \cdot a_r$, i.e.,

the splitter $s$ appears $m$ times locally. We can split $a$ into two subsequences $L = a_\ell \cdot s^x$ and $R = s^{m-x} \cdot a_r$ where we are free to choose $x \in 0..m$. We choose $x$ in such a way that $|a_\ell s^x|$ is as close to $|a|/2$ as possible. In a communication step, PEs differing in bit $j$ of their PE number exchange these sequences so that the PE with the 0-bit gets the two $L$ sequences and the PE with the 1-bit gets the $R$ sequences. These two sequences are then merged.

*Analysis:* Here we give evidence why we believe the following conjecture to be true:

*Conjecture 1:* Our robust hypercube quicksort runs in expected time

$$\mathcal{O}\left(\frac{n}{p}\log n + \beta\frac{n}{p}\log p + \alpha\log^2 p\right)$$

for arbitrary inputs with unique keys.

The proof is straight-forward for best-case inputs where all sequences are split perfectly in each step – we have $\mathcal{O}(\frac{n}{p}\log\frac{n}{p})$ time for local sorting and time $\mathcal{O}(\beta\frac{n}{p}\log p + \alpha\log p)$ for random permutation (note that this is small compared to the overall execution time for $n \ll \frac{\alpha}{\beta}\log p$). Each iteration of the main loop costs time $\mathcal{O}(\alpha\log p)$ for splitter selection and time $\mathcal{O}(\beta\frac{n}{p})$ for local data exchange and merging.

The initial random shuffling is pivotal for generalizing the result to worst case inputs. From a general point of view this is not surprising since it essentially transforms a worst case input into an average case input. A closer look shows that it simplifies the algorithm in three different respects:

The first advantage is that this makes the algorithm robust against skew also in intermediate steps. To see this, consider the path an input element travels during the execution of the algorithm. It travels from a random position (determined by the random shuffle) to its final destination (governed by the splitters used in the $\log p$ levels of recursion). **ANMERKUNG:***Element placed at random PEs within the subcube.* Hence, the average case analysis of greedy hypercube packet routing (e.g., [14]) applies that shows that overall contention is low with high probability.

The second advantage is that random shuffle also makes it possible to quickly find a good approximation of the median using the algorithm from Section III-B. To see that this also implies good overall load balance, consider the following calculation. According to Section III-B, the imbalance factor introduced in recursion level $i$ is bounded by $1+(n/2^i)^{-\gamma}$ for some positive constant $\gamma$. We use $n/2^i$ rather than the actual (already imbalanced) subproblem size here because, at the end, the PEs with highest load determine the overall imbalance. Hence, the most highly loaded subproblems are essential here. For those, $n/2^i$ is an underestimate of the subproblem size which results in an *overestimate* of the imbalance factor introduced by splitting that subproblem. Now we can estimate the overall imbalance factor $I$ as

$$\begin{aligned} I &= \prod_{i<\log p}\left(1+\left(\frac{2^i}{n}\right)^\gamma\right) = e^{\ln\prod_{i<\log p}\left(1+\left(\frac{2^i}{n}\right)^\gamma\right)} \\ &= e^{\sum_{i<\log p}\ln\left(1+\left(\frac{2^i}{n}\right)^\gamma\right)} \le e^{\sum_{i<\log p}\left(\frac{2^i}{n}\right)^\gamma} \\ &= e^{n^{-\gamma}\sum_{i<\log p}(2^\gamma)^i} \le e^{\frac{\left(\frac{p}{n}\right)^\gamma}{2^\gamma-1}} = \mathcal{O}(1) \text{ for } n \ge p \ . \end{aligned}$$

The first "$\le$" uses the Taylor series development of $\ln$ and the second one is a geometric sum. This calculation also shows that the overall number of elements in any subcube in any level of recursion is balanced. We argue that the initial random data distribution also implies a random data distribution of the elements in each subcube. Therefore, the overall balance of the number of elements also implies that the individual number of elements in each PE is not too imbalanced.[1]

Finally, random data shuffling also helps our simple local tie-breaking – each PE has a random sample of the globally available data and hence our local balancing approximates a global balancing.

## VII. Experimental Results

We present the results of a simple binomial-tree gather-merge (*GatherM*), all-gather-merge (*AllGatherM*) from Section II, robust fast work-inefficient sorting (*RFIS*) from Section V, robust quicksort on hypercubes (*RQuick*) from Section VI, robust AMS-sort (*RAMS*) from Section V, and a simple $p$-way sample sort implementation (*SSort*). We compare the results with our closest competitors, HykSort [6] (*HykSort*) and bitonic sort (*Bitonic*) which HykSort used for small inputs. We ran our experiment on JUQUEEN, an IBM BlueGene/Q based system with a 5-D torus network and spent more than four million core-hours.

We ran benchmarks with eight input instances proposed by Helman et al. [5]. We report results for the input instances *Uniform* (independent random values), *BucketSorted* (locally random, globally sorted), *DeterDupl* (only $\log p$ different keys), and *Staggered* (both designed to be difficult for hypercube-like routing). We do not show results for Gaussian, g-Group, Zero, and RandDupl as their results are more or less represented by results we show: Uniform is similar to Gaussian; BucketSorted is similar to g-Group and Zero with the difference that RQuick sorts the instances g-Group and Zero with more than $2^{17}$ elements twice as slow as RAMS; DeterDupl is similar to RandDupl. We also included the input instances *Mirrored* and *AllToOne*. The input of Mirrored on PE $i$ are random numbers between $2^{31}(m_i)/p$ and $2^{31}(m_i+1)/p$ where the bit representation of $m_i$ is the reverse bit representation of $i$. After $\frac{\log p}{2}$ recursions of a naive implementation of hypercube quicksort, $\lceil\sqrt{p}\rceil$ PEs hold $\lfloor n/\sqrt{p}\rfloor$ elements each. The first $n/p-1$ elements of AllToOne on PE $i$ are random elements from the range $[p+(p-i)(2^{32}-p)/p, p+(p-i+1)(2^{32}-p)/p]$ and the last element has the value $p-i$. At the first level of a naive implementation of $k$-way sample sort, the last $\min(p, n/p)$ PEs would send a message to PE 0.

In Section VII-A, we present results for our algorithms and our closest competitors on 262 144 cores (16 384 nodes), the largest available jobs on JUQUEEN. In Section VII-B, we

[1] A full proof may not be trivial however. For example, there *are* significant fluctuations in the individual PE loads for $n/p = \mathcal{O}(\log p)$ . However, these are likely to only introduce overhead factors dwarfed by the overhead for median selection as elements are randomly placed on the PEs associated with the subcube. We also have to take into accout that the probability of a maximal imbalance factor of $1+(n/2^i)^{-\gamma}$, introduced in recursion level $i$, decreases as the load of subcubes decreases. **TODO:***We might give a benchmark which tells us that this is not a problem we have to think about!*

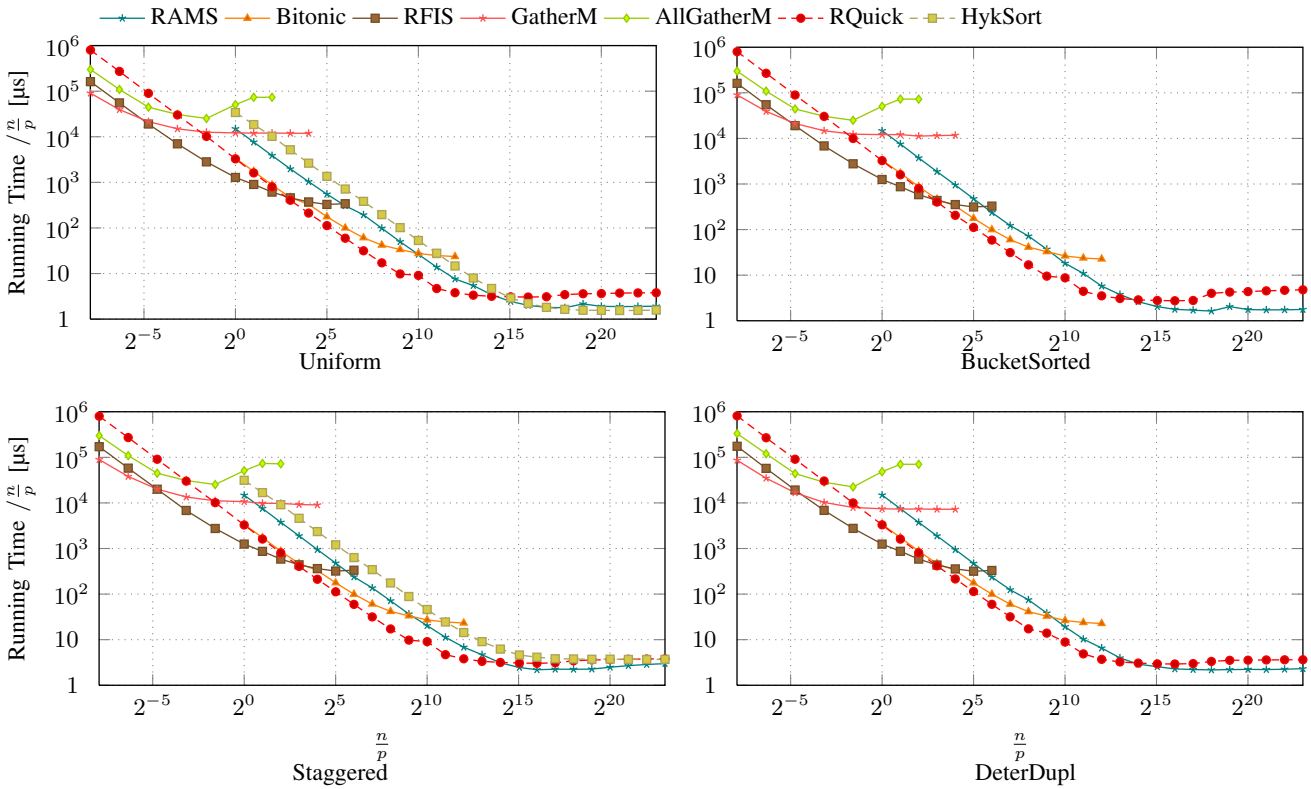


Fig. 1: Running times of each algorithm on 262 144 cores. HykSort crashes on input instances DeterDupl and BucketSorted.

compare our robust algorithms to their nonrobust versions. We show that the nonrobust counterparts are orders of magnitude slower for hard inputs. See Appendix I for more information on JUQUEEN, the input instances, and the benchmarking details. We also give implementation details in Appendix I1, describe our extensive parameter tuning on 132 072 cores in Appendix I2, and compare RAMS and HykSort to the literature in Appendix K.

### A. Input Size Analysis and Algorithm Comparison

In this section we present the results of our experiments on all algorithms executed on 262 144 cores. The smallest inputs have only one element every $3^5$ (243) PEs whereas the largest input has $2^{23}$ elements per core.

Figure 1 shows the running times of each algorithm for the most interesting input instances (Appendix J gives running time ratios of each algorithm compared to the fastest algorithm for the most interesting input instances). We now explain the most important results. (1) GatherM sorts very sparse inputs ($n/p \leq 3^{-3}$) up to 1.8 times faster than the other sorting algorithms. AllGatherM is not competitive for any input size and sorts even the sparsest input twice as slow as RFIS. However, neither fulfills the balance constraint of sorted output. (3) For all input instances, RFIS is the fastest sorting algorithm when $3^{-3} < n/p \leq 4$. For example, RFIS sorts Uniform sparse inputs up to 3.6 times faster than its competitors and a single element per core more than twice as fast as RQuick and Bitonic. (4) RQuick sorts small inputs ($2^3$ to $2^{14}$) of Uniform instances up to 3.4 times faster than any other algorithm (more than 5 times faster if we exclude our own algorithms). In general, the running times of RQuick for other input instances do not differ much for sparse and small inputs. For $2^{15}$ and $2^{16}$ elements per PE, RAMS is the fastest algorithm. For $n/p \geq 2^{17}$, the situation becomes more instance dependent. For uniform inputs, HykSort is 1.22 times faster than RAMS for the largest inputs (and 1.38 times faster for the isolated case $n/p = 2^{19}$). However, HykSort crashed on input instances DeterDupl, RandDupl, BucketSorted, g-Group, and Zero and is up to 1.7 times slower than RAMS for Staggered inputs. Overall, RAMS seems to be a good compromise between robustness and performance for large inputs. RAMS is up to 2.7 (BucketSorted) times faster than RQuick.

### B. Robustness Analysis

We now have a closer look at the impact of various algorithmic measures to improve robustness.

**Robust Quicksort:** Figure 2a depicts the running time ratio of RQuick over RQuick without redistribution (NTB-Quick). The price of robustness for simple input instances such as Uniform is an additional data redistribution. For large Uniform inputs, this slows down RQuick by a factor of up to 1.7. On the other hand, RQuick sorts skewed input instances such as Staggered and Mirrored robustly. Thus, the additional redistribution of RQuick decreases the total running time for large inputs by a factor of up to 9 ($n = 2^{15}$) and RQuick runs out of memory for $n = 2^{16}$. Our results show that RQuick sorts other input instances (BucketSorted and DeterDupl) faster

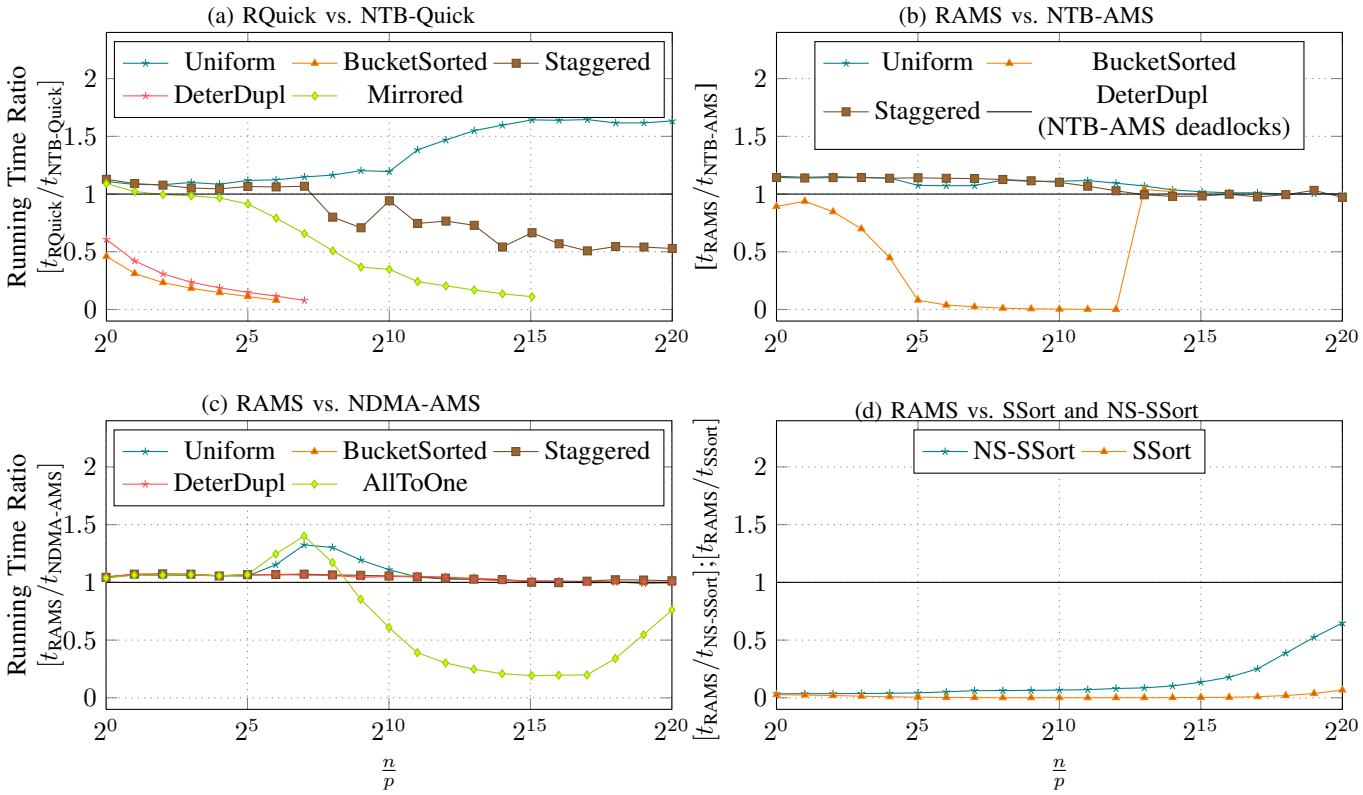


Fig. 2: Running time ratios of robust algorithms to nonrobust algorithms. 2a) Ratios of RQuick to NTB-Quick (262 144 cores). 2b) Ratios of RAMS to NTB-AMS (8 192 cores). 2c) Ratios of RAMS to NDMA-AMS (131 072 cores). 2d) Ratios of RAMS to NS-SSort (131 072 cores).

than NTB-Quick by several orders of magnitude due to tie-breaking. NTB-Quick even fails to sort such input instances when the input is large.

**Tie-Breaking of RAMS:** Figure 2b shows the running time ratio of RAMS and RAMS without tie-breaking during local data partitioning (*NTB-AMS*). When sorting small inputs without unique keys (e.g., Uniform and Staggered), the extra work to calculate splitters with tie-breaks slows down RAMS compared to NTB-AMS (factor of 1.15 for $n/p < 2^{12}$). However, this effect becomes negligible for the large inputs ($n/p \geq 2^{16}$) RAMS is intended for (i.e., when it outperforms RQuick). NTB-AMS sorts small inputs of the instance BucketSorted much slower than RAMS as small instances of BucketSorted input contain many duplicates. The non-robust version deadlocks immediately when we sort DeterDupl and Zero instances.

**Message Assignment of RAMS:** Figure 2c shows the running time ratio of RAMS using $l = 3$ levels of data redistribution and deterministic message assignment (DMA) on 131 072 cores with the same algorithm without DMA (NDMA-AMS). RAMS decides to sort Staggered, BucketSorted, and DeterDupl inputs without DMA as there would be no impact of DMA. Our experiments show that the overhead for making that decision is small. RAMS actively performs DMA for small inputs (64 to $2^{10}$) of Uniform instances leading to a certain overhead. However, for such small instances RQuick would be used preferable in any case. For the input instance AllToOne, NDMA-AMS sends $\mathcal{O}(\min(n/p, p))$ messages to PE 0 in the first level. In this case, RAMS actually can take advantage of DMA as it reduces the startups to $\mathcal{O}(k)$. This speeds up RAMS by a factor of up to 5.2. In Figure 2c, the positive effect begins for $n/p > 8k = 2^9$ elements per core and increases rapidly. The effect decreases when the time for the message exchange dominates the startup time.

**Simple Sample Sort:** The comparisons above are for recent, highly sophisticated massively parallel sorting algorithms. To give an impression how big the improvements of the new algorithms are compared to the majority of parallel sorting algorithms presented in the past for moderate number of PEs, we also compare 3-level RAMS with a simple $p$-way sample sort that delivers the data directly. Figure 2d shows that RAMS for Uniform instances is up to 1 000 times faster than SSort. RAMS is much faster even if we ignore the time for finding splitters (*NS-SSort*). When we look at the range where RAMS is faster than RQuick (at least $2^{15}$ elements per core), RAMS outperforms NS-SSort by a factor of 1.5 to 7.4. Experiments show that this effect increases as $p$ increases. Note that the curve for NS-SSort acts as a rough lower bound for any algorithm that delivers the data directly.

## VIII. Conclusion and Future Work

We show how to obtain robustness of massively parallel sorting with respect to the input size by using three different algorithms (RFIS, RQuick, and RAMS). Robustness with respect to skew can be achieved by careful randomization (RQuick) and message assignments (RAMS). Robustness with

respect to duplicate keys can be achieved at reasonable overhead using careful implicit implementations of the brute-force tie-breaking idea. We plan to release an open-source library with our algorithms.

Still further improvements are possible. We could remove the limitation to $p$ being a power of two. This will complicate the code but should not be a fundamental obstacle. For large inputs, a specialized shared-memory implementation for node-local sorting, merging and partitioning seems useful. Porting to systems with less deterministic behavior than BlueGene/Q is likely to require sophisticated implementations of collective communication and data delivery operations – the wide fluctuations of running times mentioned in [4] are already quite bad on $2^{15}$ PEs on an oversubscribed InfiniBand-Network.

*Acknowledgement.* In Appendix L we acknowledge the Gauss Centre for Supercomputing (GCS) for providing computation time, the DFG for partially funding this research and the SAP AG, Ingo Müller, Sebastian Schlag, Helman et al. [5], and Sundar et al. [6] for making their code available.

## REFERENCES

[1] M. Bader, *Space-filling curves: an introduction with applications in scientific computing*. Springer Science & Business Media, 2012, vol. 9.

[2] C. Siebert and F. Wolf, "Parallel sorting with minimal data," in *Proc. of the 18th European MPI Users' Group Meeting (EuroMPI), Santorini, Greece*, ser. Lecture Notes in Computer Science, Y. Cotronis, A. Danalis, D. Nikolopoulos, and J. Dongarra, Eds., vol. 6960. Springer, Sep. 2011, pp. 170–177.

[3] R. Cole, "Parallel merge sort," *SIAM Journal on Computing*, vol. 17, no. 4, pp. 770–785, 1988.

[4] M. Axtmann, T. Bingmann, P. Sanders, and C. Schulz, "Practical massively parallel sorting," in *Proceedings of the 27th ACM on Symposium on Parallelism in Algorithms and Architectures*, ser. SPAA '15. New York, NY, USA: ACM, 2015, pp. 13–23. [Online]. Available: http://doi.acm.org/10.1145/2755573.2755595

[5] D. R. Helman, D. A. Bader, and J. JáJá, "A randomized parallel sorting algorithm with an experimental study," *J. Parallel Distrib. Comput.*, vol. 52, no. 1, pp. 1–23, Jul. 1998.

[6] H. Sundar, D. Malhotra, and G. Biros, "Hyksort: A new variant of hypercube quicksort on distributed memory architectures," in *27th ACM Int. Conference on Supercomputing (ICS)*, 2013, pp. 293–302. [Online]. Available: http://doi.acm.org/10.1145/2464996.2465442

[7] G. E. Blelloch, C. E. Leiserson, B. M. Maggs, C. G. Plaxton, S. J. Smith, and M. Zagha, "A comparison of sorting algorithms for the connection machine CM-2." in *3rd Symposium on Parallel Algorithms and Architectures*, 1991, pp. 3–16.

[8] P. J. Varman, S. D. Scheufler, B. R. Iyer, and G. R. Ricard, "Merging multiple lists on hierarchical-memory multiprocessors," *J. Par. & Dist. Comp.*, vol. 12, no. 2, pp. 171–177, 1991.

[9] A. Gerbessiotis and L. Valiant, "Direct bulk-synchronous parallel algorithms," *Journal of Parallel and Distributed Computing*, vol. 22, no. 2, pp. 251–267, 1994.

[10] J. Singler, P. Sanders, and F. Putze, "MCSTL: The multi-core standard template library," in *13th Euro-Par*, ser. LNCS, vol. 4641. Springer, 2007, pp. 682–694.

[11] K. E. Batcher, "Sorting networks and their applications," in *AFIPS Spring Joint Computing Conference*, 1968, pp. 307–314.

[12] S. L. Johnsson, "Combining parallel and sequential sorting on a boolean n-cube," in *Proceedings of the International Conference on Parallel Processing*, 1984, pp. 444–448.

[13] W. Gropp, E. Lusk, and A. Skjellum, *Using MPI*. MIT Press, 1995.

[14] T. Leighton, *Introduction to Parallel Algorithms and Architectures*. Morgan Kaufmann, 1992.

[15] P. Sanders and T. Hansch, "Efficient massively parallel quicksort," in *Solving Irregularly Structured Problems in Parallel*. Springer, 1997, pp. 13–24.

[16] B. C. Dean, R. Jalasutram, and C. Waters, "Lightweight approximate selection," in *Algorithms-ESA 2014*. Springer, 2014, pp. 309–320.

[17] B. Wagar, "Hyperquicksort: A fast sorting algorithm for hypercubes," *Hypercube Multiprocessors*, vol. 1987, pp. 292–299, 1987.

[18] Y. Lan and M. A. Mohamed, "Parallel quicksort in hypercubes," in *Proceedings of the 1992 ACM/SIGAPP Symposium on Applied Computing: Technological Challenges of the 1990's*, ser. SAC '92. New York, NY, USA: ACM, 1992, pp. 740–746.

[19] E. Solomonik and L. Kale, "Highly scalable parallel sorting," in *IEEE International Symposium on Parallel Distributed Processing (IPDPS)*, April 2010, pp. 1–12.

[20] P. Sanders and T. Hansch, "On the efficient implementation of massively parallel quicksort," in *4th International Symposium on Solving Irregularly Structured Problems in Parallel*, ser. LNCS, G. Bilardi, A. Ferreira, R. Lüling, and J. Rolim, Eds., no. 1253. Springer, 1997, pp. 13–24.

[21] L. Valiant, "A bridging model for parallel computation," *Communications of the ACM*, vol. 33, no. 8, pp. 103–111, 1994.

[22] R. Cole and V. Ramachandran, "Resource oblivious sorting on multicores," in *Automata, Languages and Programming (ICALP)*, ser. LNCS. Springer, 2010, vol. 6198, pp. 226–237.

[23] G. E. Blelloch, P. B. Gibbons, and H. V. Simhadri, "Low depth cache-oblivious algorithms," in *22nd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 2010, pp. 189–199.

[24] M. Ikkert, T. Kieritz, and P. Sanders, "Parallele Algorithmen," October 2009, course notes.

[25] P. Sanders, J. Speck, and J. L. Träff, "Two-tree algorithms for full bandwidth broadcast, reduction and scan," *Parallel Computing*, vol. 35, no. 12, pp. 581–594, 2009.

[26] P. Sanders and S. Winkel, "Super scalar sample sort," in *12th European Symposium on Algorithms (ESA)*, ser. LNCS, vol. 3221. Springer, 2004, pp. 784–796.

[27] T. Hoefler, C. Siebert, and A. Lumsdaine, "Scalable communication protocols for dynamic sparse data exchange," *SIGPLAN Not.*, vol. 45, no. 5, pp. 159–168, Jan. 2010.

[28] Y. Feng, M. Straka, T. di Matteo, and R. Croft, "MP-sort: Sorting at scale on blue waters," 2014, https://cug.org/proceedings/cug2015_proceedings/includes/files/pap179.pdf accessed Jan 17, 2015.

[29] J. Wang, Y. Wu, H. C. Zhipeng Tang, Z. Lv, B. Lu, Y. Tao, C. L. Jingren Zhou, and H. Tang, "Fuxisort," 2015, http://sortbenchmark.org/FuxiSort2015.pdf accessed Apr. 9, 2016.

[30] Z. Zhang, C. Li, Y. Tao, R. Yang, H. Tang, and J. Xu, "Fuxi: a fault-tolerant resource management and job scheduling system at internet scale," *Proceedings of the VLDB Endowment*, vol. 7, no. 13, pp. 1393–1404, 2014.

[31] M. Stephan and J. Docter, "Jülich Supercomputing Centre. JUQUEEN: IBM Blue Gene/Q Supercomputer System at the Jülich Supercomputing Centre," *Journal of large-scale research facilities*, 2015, 1, A1. http://dx.doi.org/10.17815/jlsrf-1-18.

[32] P. Sanders, S. Schlag, and I. Müller, "Communication efficient algorithms for fundamental big data problems," in *IEEE Int. Conf. on Big Data*, 2013.

## APPENDIX

### A. Model of Computation

A common abstraction of communication in supercomputers is the (symmetric) single-ported message passing model. It takes time $\alpha + l\beta$ to send a message of size $l$ machine words. The parameter $\alpha$ defines the startup overhead of the communication. The parameter $\beta$ defines the time to communicate a machine word. For simplicity, we assume that the size of a machine word is equivalent to the size of a data element. For example, broadcast, reduction, and prefix sums can be implemented to run in time $\mathcal{O}(\beta l + \alpha \log p)$ [11], [25] for vectors of size $l$. We have $\alpha \gg \beta \gg 1$ where our unit is the time for executing a simple machine instruction. Most of the time, we treat $\alpha$ and $\beta$ as variables in our asymptotic analysis in order to expose effects of latency and bandwidth limitations.

### B. Hypercube Algorithms

**Algorithm 3** Hypercube algorithm design pattern

```
Computation on PE i
for j:= d − 1 downto 0 do                    ▷ or 0..d − 1
    send some message m to PE i ⊕ 2^j
    receive message m′ from PE i ⊕ 2^j
    perform some local computation using m and m′
```

A hypercube network of dimension $d$ consists of $p = 2^d$ PEs numbered $\{0, \ldots, p-1\}$. Two nodes $a$ and $b$ are connected along dimension $i$ if $a = b \oplus 2^i$. For this paper, hypercubes are not primarily important as an actual network architecture. Rather, we extensively use communication in a conceptual hypercube as a design pattern for algorithms. More specifically the hypercube Algorithms 3 iterate through the dimension of the hypercube. Depending on how this template is instantiated, one achieves a large spectrum of global effects. For example, by repeatedly summing an initial local value $a$, one gets an *all-reduce*, i.e., all PEs get the global sum of the local values in time $\mathcal{O}((\alpha + \beta|a|)\log p)$. Similarly, if we replace addition by concatenation, we perform an *all-gather* operation (i.e., all PEs get all the local values) which runs in time $\mathcal{O}(\beta p|a| + \alpha \log p)$. If the $a$s are a sorted sequences and we replace concatenation by merging, all PEs get the elements in all the local $a$s in sorted order using time $\mathcal{O}(\beta p|a| + \alpha \log p)$. We call this operation *all-gather-merge*.

We can also use a hypercube algorithm for routing data – in iteration $j$, a data object destined for PE $t$ and currently located on PE $i$ is moved if $t$ and $i$ differ in bit $j$ (e.g. [14]). This has the advantage that we need only $\mathcal{O}(\log p)$ startup overheads overall. However, for worst case inputs, even if every PE sends and receives only a single object, the required time can be $\Omega(\sqrt{p})$ since many objects can visit the same intermediate nodes. However, it is known that for random start or destination nodes, the running time remains $\mathcal{O}(\log p)$ [14].

To describe and understand hypercube algorithms, we need the concept of a subcube. A $j$-dimensional subcube consists of those PEs whose numbers have the same bits $j..d-1$ in their binary representation.

### C. Randomized Shuffling on Hypercubes

It is a folklore observation (e.g. [15]) that data skew can be removed by shuffling the input data randomly. This is actually implemented for the sample sort algorithm of Helman et al. [5] by directly sending each element to a random destination. Note that this incurs an overhead of about $\alpha p + \beta n/p$. We propose to use this technique for small inputs where we need smaller latency. This can be achieved by routing the data according to a hypercube communication pattern – sending each element to a random side in each communication step. To achieve even slightly better load balance, we actually split local data in two random halves in each communication step. The resulting running time is

$$\mathcal{O}(\left(\alpha + \beta\frac{n}{p}\right)\log p) \ .$$

A naive approach labels each element with a random destination. Then, a hypercube algorithm redistributes the data. This approach increases the communication volume by a factor of two.

### D. Sorting Algorithms from Tiny to Huge Inputs

*1) Fast Work-Inefficient Ranking:* We recently proposed a simple but fast ranking algorithm for small inputs [4]. The PEs are arranged in an array of size $\mathcal{O}(\sqrt{p}) \times \mathcal{O}(\sqrt{p})$ and each PE has $\mathcal{O}(\frac{n}{p})$ elements. We first sort the elements locally in $\mathcal{O}(\frac{n}{p}\log\frac{n}{p})$. Then we all-gather-merge the elements of PEs in the same row as well as between PEs in the same column in $\mathcal{O}(\alpha \log p + \beta n/\sqrt{p})$. Afterwards, each PE stores the elements of all PEs in its row and column respectively in a sorted way. Then the PEs rank each element received from their column in elements received from their row. Finally, we sum up the local ranks in each column with an allreduce operation. The result is that each PE knows the global rank of all inputs elements in its row. In theory this is a very good algorithm for $n = \mathcal{O}(\sqrt{p})$ since we get only logarithmic delay. However, since all other algorithms need $\Omega(\log^2 p)$ startup overheads, Fast work-inefficient ranking is interesting up to $n = \mathcal{O}(\frac{\alpha}{\beta}\sqrt{p}\log^2 p)$. In Subsection F we make the algorithm robust against identical keys and show how to convert its output to a classical sorted permutation of the input.

*2) Bitonic Sort:* Bitonic sort [11], [12] first sorts locally and then performs $\mathcal{O}(\log^2 p)$ pairwise exchange and merge operations.

For small inputs this running time is dominated by $\log^2 p$ startup overheads. This gives it a similar running time as the parallel quicksort algorithms to be discussed next. However, for $n = \omega(p\alpha/\beta)$ the term $\beta\frac{n}{p}\log^2 p$ dominates – all the data is exchanged $\log^2 n$ times. This makes it unattractive for large inputs compared to quicksort and other algorithms designed for large inputs. Indeed, only for $n$ superpolynomial in $p$ ($n = \Omega(p^{\log p})$), bitonic sort eventually becomes efficient. But for such large inputs algorithms like sample sort are much better since they exchange the data only once.

*3) Parallel Quicksort:* Since quicksort is one of the most popular sorting algorithms, it is not surprising that there are also many parallel variants. For distributed memory, a variant using the hypercube communication pattern is attractive since it is simple and can exploit locality in the network. A recursive subproblem is solved by a subcube of the hypercube. The splitter is broadcast to all PEs. Elements smaller than the pivot are moved to the 0-subcube and elements larger than the pivot are moved to the right 1-subcube. Since this hypercube quicksort obliviously splits the PEs in half, the imbalance accumulating over all the levels of recursion is a crucial problem. In its simplest original form by Wagar [17], PE 0 uses its local median as a pivot. This is certainly not robust against skew and even for average case inputs it only works for rather large inputs. One can make this more robust – even in a deterministic sense – by using the global median of the local medians [18]. However, this introduces an additional $\beta p$ term into the communication complexity and thus defeats the objective of having polylogarithmic execution time. Therefore in Subsection VI we propose to use a median selection algorithm that is both fast and accurate.

Non-hypercube distributed memory quicksort is also an interesting option where we can adapt the number of PEs working on a recursive subproblem to its size. However, it leads to more irregular communication patterns and has its own load balancing problems because we cannot divide PEs fractionally. Siebert and Wolf [2] exploit the special case where this is no problem, when $n = p$.

*4) Generalizing Quicksort:* For large inputs it is a disadvantage that quicksort exchanges the data at least $\log p$ times. We can improve that by partitioning the input with respect to $k-1$ pivots at once. This decreases the number of times data is moved to $\mathcal{O}(\log_k p)$. However, the price we pay is latency $\Omega(\alpha k)$ for delivering data from $k$ partitions to $k$ different PEs. This gives us a lower bound of

$$\Omega\left(\frac{n}{p}\log n\right) + \beta\frac{n}{p}\log_k p + \alpha k \log_k p$$

for running generalized quicksort. Getting an algorithm with a matching upper bound is not easy though. We have to find pivots that partition the input in a balanced way and we have to execute the more complex data delivery problems efficiently.

Many algorithms have been devised for the special case $k = p$. Sample sort [7] is perhaps most well known because it is simple and effective. It achieves the above bound if a sample of size $S = \Omega(p \log n)$ is sorted using a parallel algorithm. Then, every $S/k$-th sample is used as a splitter. This algorithm achieves the desired bound for $n = \mathcal{O}(p^2/\log p)$. Solomonik and Kale [19] describe the best scaling single level algorithm we have seen and scales to $2^{15}$ PEs for large inputs. One can even achieve perfect partitioning by finding optimal splitters in a multiway mergesort approach [4], [8], [10]. This requires $n = \mathcal{O}(p^2 \log p)$ for achieving the desired bound.

Gerbessiotis and Valiant [9] develop a multilevel sample sort for the BSP model [21]. However, implementing the data exchange primitive of BSP requires $p$ startup overheads in each round. See [4] for a more detailed discussion of previous theoretical multilevel algorithms.

Sundar et al. develop and implement *HykSort*, a generalized hypercube quicksort [6]. HykSort uses a sophisticated recursive sampling based algorithm for finding high quality approximate splitters that is something like a compromise between the single shot algorithms used in sample sort and the exact (but slower) algorithms used for multiway mergesort. This makes HykSort "almost" robust against skew. However, in the worst case, data may become severely imbalanced. The paper mentions measures for solving this problem, but does not analyze them. Furthermore, HykSort uses the operations `MPI_Comm_Split` whose current implementations need time $\Omega(\beta p)$. This is why we put a "$\geq$" in the corresponding row of Table I. HykSort is also not robust with respect to duplicate keys.

*AMS-sort* [4] our recently published $k$-way sample sort algorithm which guarantees a maximum imbalance of $\epsilon\frac{n}{p}$ for arbitrary inputs. Within each recursion, AMS-sort sorts $\mathcal{O}(k \log k)$ samples in parallel. Then it selects and distributes $b \cdot k$ splitters where $b = \mathcal{O}(\frac{1}{\sqrt{1+\epsilon}-1})$. After the local data has been partitioned into $b \cdot k$ partitions, a greedy algorithm assigns global partitions to PE subcubes by minimizing the load imbalance between subcubes. Then each PE calculates for its partitions the target PEs within the subcubes. This ensures perfect load balancing within target subcubes. Thus, by going away from a (even generalized) hypercube communication pattern, AMS-sort avoids the data imbalance that may make HykSort inefficient. Our paper [4] describes (but does not implement) a way to select communication partners such that each PE sends and receives $\Theta(k)$ messages for a $k$-way data exchange. Similarly, a tie-breaking technique is proposed (but not implemented) which simulates unique keys with very low overhead. In Subsection G we describe our robust implementation of AMS-sort.

### E. More Related Work

This paper was highly influenced by two landmark papers. 25 years ago, Blelloch et al. [7] carefully studied a large number of algorithms, selected three of them (bitonic, radix, and sample sort), and compared them experimentally on up to $2^{15}$ PEs. Most other papers are much more focused on one or two (more or less) new algorithms sometimes missing the big picture. We felt that it was time again for a more wide view in particular because it became evident to us that a single algorithm is not enough on the largest machine and varying $n/p$. Helman et al. [5] – concentrating on sample sort – took robustness seriously and made systematic experiments with a wide variety on input distributions which form the basis for our experiments. They also propose initial random shuffling to make samplesort robust against skewed input distributions (this was already proposed in [20]).

### F. Robust Fast Work-Inefficient Sorting

We propose a fast sorting algorithm (RFIS) for sparse and very small inputs that is robust against duplicate keys. Our algorithm computes the rank of each element in the global data. Then, optionally, a data delivery routine sends the elements to the appropriate PEs according to their ranks. For ranking

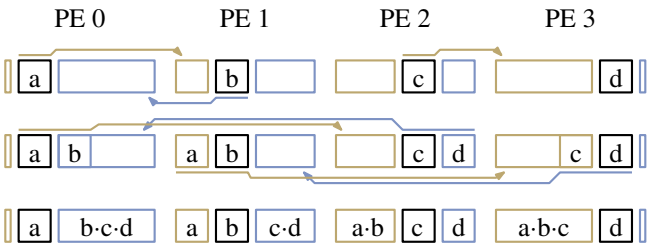


Fig. 3: All-gather-merge keeps track where elements came from. Each PE has three buckets; one bucket stores elements which came from the left (brown), one bucket stores elements which came from the right (blue), and one bucket stores elements initially provided by this PE (black).

the input data, we use the algorithm we recently proposed [4] (see also [24] and D1). Unfortunately, the proposed approach calculates the same rank for equal elements. However, our data delivery techniques require unique ranks in the range of $0..n-1$ as input to achieve balanced output.

Consider a tie-breaking policy that logically treats an element as a quadruple $(x, r, c, i)$ where $x$ is the key, $r$ the row, $c$ the column, and $i$ the position in the locally sorted data. We can then define a total ordering between these quadruples as their lexicographical order. A modified all-gather-merge and a modified compare function implement that policy without communication overhead for communicating the $(r, c, i)$ information. For the all-gather-merge in a row, each PE uses three arrays $\{\leftarrow, H, \rightarrow\}$ to keep track which elements came from the left ($\leftarrow$), which elements came from the right ($\rightarrow$) and which elements are local elements ($H$). Assume that we implement all-gather-merge with a hypercube algorithm starting from the lowest dimension. When a PE receives data from left (right), *all* elements in that message come from left (right) and are merged into the array $\leftarrow$ ($\rightarrow$). Analogously, for the all-gather-merge in a column, PEs merge elements from above into the array $\uparrow$ and elements from below into the array $\downarrow$. This classification suffices to perform tie-breaking according to the lexicographic ordering of the quadruples. Figure 3 shows the all-gather-merge algorithm in a row. We now update the compare function on PE $(r, c)$ between an element $a = (x, r, C, i)$ from row $r$ and an element $b = (y, R, c, j)$ from column $c$ where the arrays $C \in \{\leftarrow, H, \rightarrow\}$ and $R \in \{\uparrow, H, \downarrow\}$ are used as implicit element labels. Furthermore, $i$ ($j$) specifies the position among the locally sorted data if $C = c$ ($R = r$) and is undefined otherwise. Then the following table defines the updated compare function of comparing $a > b$

| $b \setminus a$ | $\leftarrow$ | $c$ | $\rightarrow$ |
|---|---|---|---|
| $\uparrow$ | $a \geq b$ | $a \geq b$ | $a \geq b$ |
| $r$ | $a > b$ | $i > j$ | $a \geq b$ |
| $\downarrow$ | $a > b$ | $a > b$ | $a > b$ |

When we rank row elements into column elements, we choose the correct compare function just once for each combination of row array $\{\leftarrow, H, \rightarrow\}$ and column array $\{\uparrow, H, \downarrow\}$.

After we computed the ranks of the input elements, we redistribute them to get sorted output.

### G. Robust Multi-Level Sample Sort

The prototypical implementation of AMS in [4] does not ensure that each PE communicates with $\mathcal{O}(k)$ PEs in each level of recursion. Similarly, [4] proposes but does not implement a tie-breaking scheme to become robust against duplicate keys. Here, we present a complete implementation of AMS-sort (RAMS) including both improvements outlined above.

Even for worst case inputs, the following partitioning approach splits input as balanced as for input with unique keys. To do so, we first pick random samples of the local data and add their positions in the input as tie-breakers to the samples. Then we rank the samples with the fast work-inefficient ranking algorithm (see Appendix D1). Next, we select splitters from the samples and distribute them to all PEs. Finally, the partitioner of Super Scalar Sample Sort [26] assigns all local elements to a partition. We modify the element classifier of Super Scalar Sample Sort to support tie-breaking. Initially, we classify each input element using its original key. If this key is equal to the original key of the bounding splitter, we repeat the search using the positions of the splitters as tie-breakers. To become cache efficient, we perform multiple partitioning passes if the number of partitions is large.

Our robust version of AMS-sort implements the *deterministic message assignment* proposed for AMS-sort. The message assignment algorithm first routes information about messages to groups of PEs where the actual message belongs. Then each group calculates for each message one or more addresses. Finally, these addresses are sent back to the PEs which will send the actual messages back to the PE group. However, the PEs do not know the number of assigned addresses and thus do not know the number of messages they have to receive. The algorithm NBX [27], a dynamic sparse data exchange technique, uses an non-blocking barrier to identify when the message exchange is complete. We use this method to efficiently exchange the addressees in $\mathcal{O}(\alpha \log p + \alpha k)$ time.

### H. Experimental Comparison of Median Approximation Trees

We experimentally evaluate the median approximations of the binary-tree selection proposed in Subsection III-B and the ternary-tree selection proposed by Dean et al. [16]. We show that the binary-tree selection gives better median approximations than the ternary-tree selection and the rank can also be bounded by $n/2(1 \pm cn^{-\gamma})$ with better constants $c$ and $\gamma$. Our benchmark runs both selection algorithms on uniformly distributed random integers in the range $[0, 2^{32} - 1]$ of up to $2^{20}$ elements. The input into the binary-tree selection is a power of two whereas the input into the ternary-tree selection is a power of three. We execute both algorithms 2 000 times for each input size. We measure the rank $r$ for each output element and calculate the rank error

$$\left| \frac{r}{n-1} - \frac{1}{2} \right| .$$

Figure 4a shows the rank error in worst cast. Compared to the ternary-tree selection, the maximal rank error over 2 000 runs of the binary-tree selection is smaller for all input sizes. Moreover the maximal rank error of the binary-tree selection can tightly bounded by $1.44n^{-0.39}$ but the maximal rank

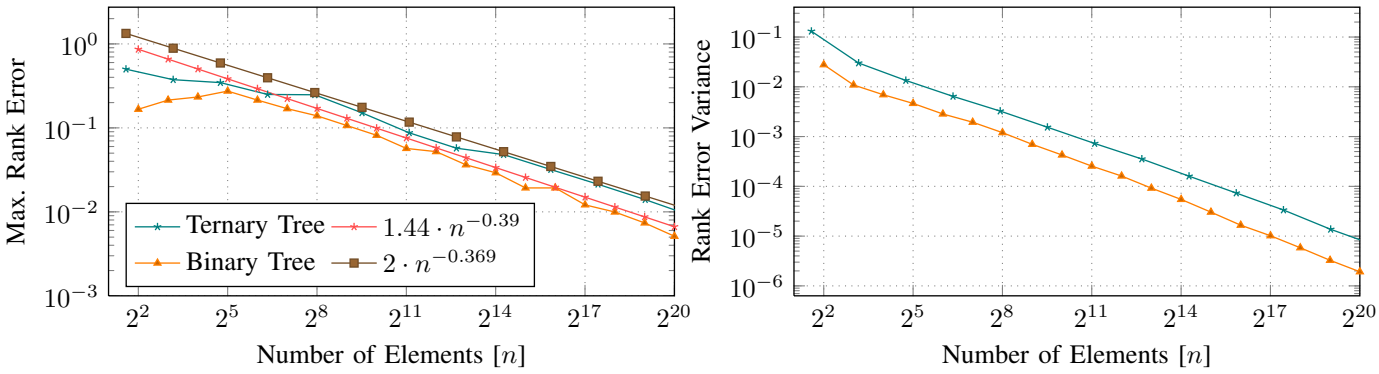


(a) (b)

Fig. 4: Maximal rank error (a) and variance of the rank error (b) of the ternary-tree median approximation and the binary-tree median approximation.

error of the ternary-tree selection is just limited by $2n^{-0.369}$. Figure 4b depicts the variance of the rank error for the binary-tree selection and the ternary-tree selection. For small input sizes, the variance of the binary-tree median approximation is smaller by a factor of two. For large inputs, the difference in the variance even increase to a factor of three.

### I. Experimental Results

We ran our experiment on JUQUEEN, an IBM BlueGene/Q based system of 56 midplanes, each with 512 compute nodes. The maximum number of midplanes available to us was 32. Each compute node has one IBM PowerPC A2 16-core processor with a nominal frequency of 1.6 GHz and 16 GByte of memory. A 5-D torus network with 40 GB/s and 2.5 µsec worst case latency connects the nodes. The user can specify the number of midplanes per torus dimension. In our experiments, we configured the tori such that the size of the dimensions are as equal as possible.

We ran benchmarks with nine input instances of 64-bit floating point elements. We use the input instances proposed by Helman et al. [5]: *Uniform* (independent random values), *Gaussian* (independent random values with Gaussian distribution), *BucketSorted* (locally random, globally sorted), *DeterDupl* (only $\log p$ different keys), *RandDupl* (32 local buckets of random size, each filled with an arbitrary value from 0..31), *Zero* (all elements are equal), *g-Group* with $g := \sqrt{p}$, and *Staggered* (both designed to be difficult for hypercube-like routing). We also included the input instances *Reverse* (reverse sorted input) and *AllToOne* where the first $n/p-1$ elements on PE $i$ are random elements from the range $[p+(p-i)(2^{32}-p)/p, p+(p-i+1)(2^{32}-p)/p]$ and the last element has the value $p-i$. At the first level of a naive implementation of $k$-way sample sort, the last $\min(p, n/p)$ PEs would send a message to PE 0.

We repeat each measurement six times, ignore the first run (used as a warmup) and average over the remaining five runs. Overall, we spent more than four million core-hours. We do not show error bars since the ratio between the maximum execution time and the average execution time is consistently less than a few percent. The only exception is all-gather-merge which sometimes fluctuates by up to 50 % for $n = p$.[2] The time for building MPI communicators (which can be considerable) is not included in the running time since this can be viewed as a precomputation that can be reused over many runs with arbitrary inputs.

*1) Implementation Details:* We call inputs *dense* when $n/p \geq 1$ and sparse when $n/p < 1$. *Sparsity-factor* $i$ means that only every $i$-th PE holds an input elements. The implementations GatherM, AllGatherM, RFIS, RQuick, and RAMS are in general capable of sorting all inputs. Exceptions are Bitonic which fails to sort sparse inputs and HykSort which fails to sort inputs with many duplicates.

For RFIS, we proposed in Subsection V a hypercube exchange operation to route the elements to their target PEs. In contrast, our implementation delivers elements in single messages which is more efficient for $n/p < \log p$. For those input sizes, we expect RFIS to be faster than its competitors. Axtmann el al. [4] use a scattered all-reduce operation to reduce the local ranks in their variant of FIS. We use `MPI_Allreduce` instead without observing significant overheads due to this decision.

$l$-level RAMS selects $b = \frac{2}{\sqrt[l]{1+\epsilon}-1}$ splitters to guarantee a maximum imbalance of $\epsilon = 0.2$. In our experiments, the imbalance was always smaller than 0.1 (except for $n/p \leq 16$). SSort uses a sample size of $16 \log p$ elements per PE and calls `MPI_Alltoallv` to exchange data. The implementation uses the `MPI_Datatype MPI_Byte` for every `C++` data type. When we used other `MPI_Datatypes`, the `MPI_Alltoallv` call ran into a machine specific deadlock. Sundar et al. [6] implemented a nonrobust and a robust version of $k$-way sample sort which both use `MPI_Alltoall` and `MPI_Alltoallv`. We did not include their implementation in our benchmarks as these MPI-calls deadlocked. The deadlocks may also be a machine specific problem as the implementation successfully ran on our institute server. Siebert and Wolf [2] developed *Minisort* for the case $n = p$. Unfortunately, the exact source code used in their experiments is not available any more. The algorithms are written in `C++11` and compiled

[2] An interesting observation is that the deterministic algorithm Bitonic has only negligible fluctuations. This indicates that the fluctuations introduced by the machine itself are very small.

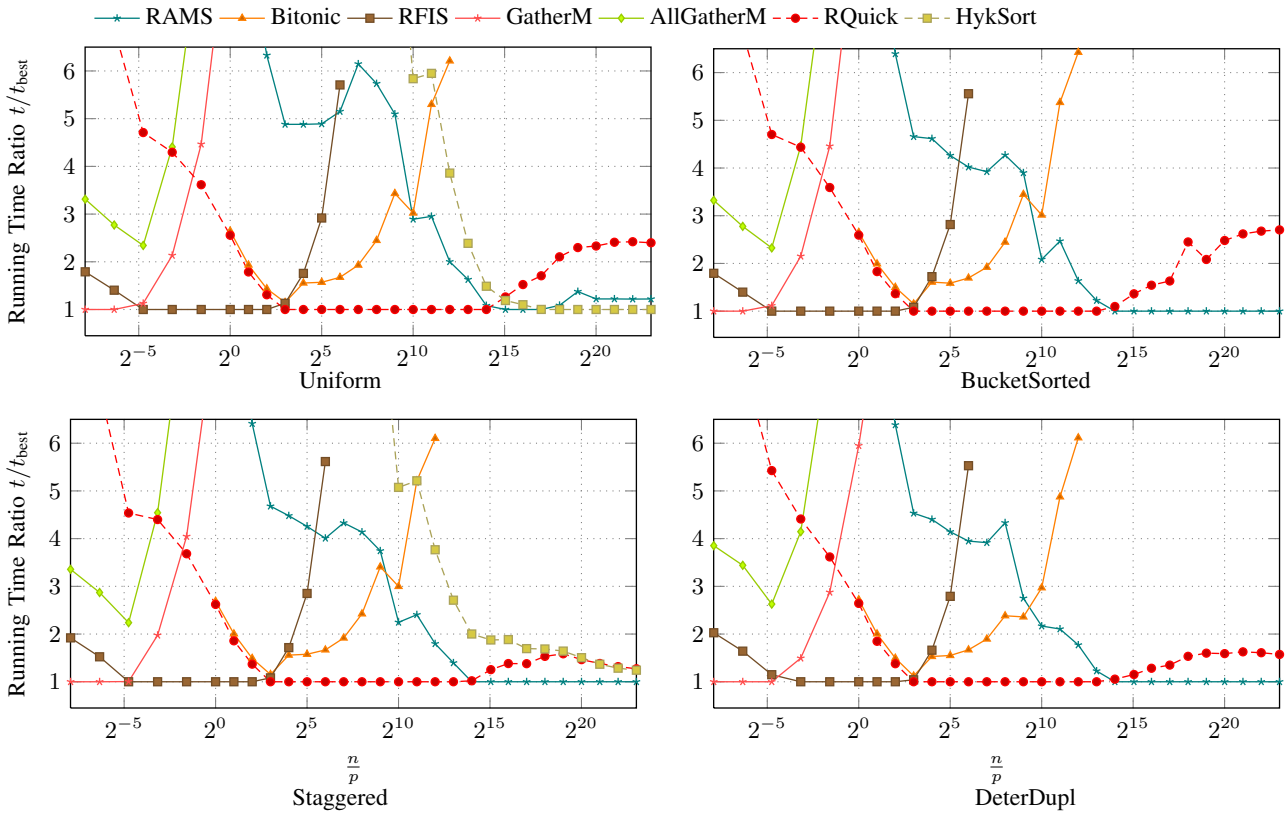


Fig. 5: Running ratios of each algorithm to the fastest algorithm on 262 144 cores. HykSort crashes on input instances DeterDupl and BucketSorted.

with version 4.8.1 of the GNU compiler collection, using the optimization flag *-O2*. For inter-process communication, we used the IBM `mpich2` library version 1.5, Argonne.

*2) Parameter Tuning:* We performed extensive parameter tuning on 131 072 cores, half of the available number due to a limited number of available core hours. Each algorithm has a range of input sizes where it is designed to be the fastest. We optimized each algorithm to run as fast as possible in this input range. Sundar et al. [6] parallelized the initial sorting step of Bitonic. In our experiments, we use Bitonic with one thread per process since the merging step is not parallelized. Just for $n = p$, using more threads per process is slightly faster than Bitonic with one thread per process. We tested HykSort with all combinations of $k = 16, 32, 64$ and $2, 8, 16$ threads per process and sorted Uniform inputs of up to $2^{20}$ elements. HykSort did not terminate when we executed the algorithm with one thread per process. HykSort with $k = 32$ and 8 threads is the fastest algorithm overall (or is as fast as RAMS) for at least $2^{17}$ elements. RAMS with three levels is the fastest algorithm overall for input of size $2^{15}$ and $2^{16}$. More levels speed up RAMS for small inputs (up to 50%) and less levels slightly speed up RAMS for larger inputs.

### J. Algorithm Comparison with Running Time Rations

Figure 5 shows the running time ratios of each algorithm compared to the fastest algorithm for the most interesting input instances.

### K. Comparison to the Literature

We compare our results to the literature and the web. MP-sort [28] is a single-level multiway mergesort running on up to 160 000 cores of a Cray XE6. For 128 GB of data it is 3 635 times slower than ours for $p = 262\,144$ and 137 GB. The 2015 winner of the Sort Benchmark (sortbenchmark.org), FuxiSort [29], [30] sorts 7.7 TB of 100 byte elements with 10 byte random keys using 41 496 cores in 58s (MinuteSort). That is 9 times longer than RAMS and HykSort take for sorting 1.1 times as much data and 15.7 times more (8-byte) elements using 6.3 times more cores. FuxiSort has to read/write from/to disks. However, the disks of that system are not the bottleneck.

### L. Acknowledgment

The authors gratefully acknowledge the Gauss Centre for Supercomputing (GCS) for providing computing time through the John von Neumann Institute for Computing (NIC) on the GCS share of the supercomputer JUQUEEN [31] at Jülich Supercomputing Centre (JSC). GCS is the alliance of the three national supercomputing centres HLRS (Universität Stuttgart), JSC (Forschungszentrum Jülich), and LRZ (Bayerische Akademie der Wissenschaften), funded by the German Federal Ministry of Education and Research (BMBF) and the German State Ministries for Research of Baden-Württemberg (MWK), Bayern (StMWFK) and Nordrhein-Westfalen (MIWF). Special thanks go to SAP AG, Ingo

Müller, and Sebastian Schlag for making their 1-factor algorithm [32] available. We gratefully thank Helman et al. [5] which provided generators for seven input instances and Sundar et al. [6] for making HykSort available. This research was partially supported by DFG project SA 933/11-1.